



22883

PATENT TRADEMARK OFFICE

155.1006.01

This application is filed in the name of the following inventor(s):

<i>Inventor Name</i>	<i>Citizenship</i>	<i>Residence City and State</i>
Stephane KASRIEL	France	Mountain View, California
Walter MANN	United States	San Francisco, California

The assignee is *Fireclick, Inc.*, a corporation having a place of business at 289 South San Antonio Road, Suite 150, Los Altos, California 94022.

Title of the Invention

Efficient Compression Using Differential Caching

Background of the Invention

1. *Field of the Invention*

This invention relates to reducing the bandwidth and computing resources required to transmit a web page over a network.

2. *Related Art*

It is desirable to transmit web pages in such a way as to minimize the bandwidth and other computing resources required to transmit the web page from a server device to a client device.

One popular technique for minimizing bandwidth requirements and computing resources is to store the web pages at multiple mirroring servers throughout the network. By “pushing content to the edge of a network” (that is, caching the content and serving it from a mirroring server), it is possible to minimize the distance (that is, geographical distance or as measured by network topology) that information must travel before reaching a destination client device. Although redistribution of the load among multiple mirroring servers improves the quality of service by minimizing distance, it does not necessarily affect the bit rate at which a web page is transmitted.

Another technique for minimizing bandwidth requirements involves compression of the web page. Compression generally involves the use of a computer program such as gzip, glib or some other similar program. Since HTML is very compressible (for example, the use of glib can result in 73.8 % compression of HTML), there is a significant minimization in the amount of bandwidth required for transmission of web pages. One drawback to this technique however, is that compressing each web page is a relatively inefficient use of computing resources. Every time a web page is sent to a user,

1 the page must be compressed, regardless whether the page was compressed in the past.
2 For example, every time an organizational home page is sent from a server to a client de-
3 vice, computing resources are devoted to compression of that home page. In a second
4 example, every time a server sends similar pages (such two different instances of an order
5 form from an on-line retailer), those pages must be separately compressed without regard
6 for their similarity.

7
8 Accordingly, it would be desirable to provide a technique for serving rela-
9 tively non-static content for delivery in a content delivery network.

10 11 Summary of the Invention

12
13 In a first aspect of the invention, a template for a web page or for a set of
14 web pages is identified and compressed at an originating server. A template can include
15 either (1) an entire web page or (2) those elements in a web page that are relatively un-
16 changing. Examples of templates include:

- 17
- 18 • Entire web pages
 - 19 • Order form pages without the personalized information
 - 20 • Stock update pages without the ticker information
 - 21 Product specs without the product descriptors
 - 22 Ticket sale pages, without information about a particular event

1 Portions of comparable web pages wherein a portion of the web page does not
2 change frequently.

3
4 A library of these compressed templates is stored at the originating server and the mir-
5 roring servers so as to make it available for subsequent requests. Generally, compression
6 of the template only occurs once. The template does not need to be compressed again,
7 unless changes to the template itself are made (for example, changes to the style of the
8 template). In this respect, the invention is very different from the prior art of HTML
9 compression which requires compression of a web page every time the page is sent. By
10 minimizing the number of times that compression takes place, fewer computational re-
11 sources need be allocated toward compression operations. In a preferred embodiment,
12 the compressed template is cached and served from a remote mirroring server. In other
13 embodiments, the compressed template is cached and served from the originating server.
14 When a client device requests a web page, the compressed template is sent from either
15 the mirroring server (if the template is present) or from the originating server (if the tem-
16 plate is not present at the mirroring server) by way of the mirroring server to the client
17 device.

18
19 In a second aspect of the invention, delta information from a set of web
20 pages is identified and compressed at the originating server. Unlike templates which in-
21 clude substantial similarities, delta information refers to a set of differences such as dif-
22 ferences between instances of a particular web page. For example, in a web page of a

completed order form, the web page may be divided into (1) a template, including the un-
changing features of the page (or the entire page if there is no dynamic information in the
page); and (2) delta information comprising specific personalized information pertaining
to an individual order. In a preferred embodiment, the delta information is compressed at
the originating server and sent from the originating server to the client. In a preferred
embodiment, delta information is compressed each time it is transmitted. Given that the
delta information is usually relatively small, the amount of computational resources de-
voted to its compression is also relatively small. This selective compression minimizes
the bandwidth needed to transmit the page while simultaneously minimizing the com-
puting resources allocated for compression. In other embodiments, the delta information
may be compressed once and cached until such time that the compressed delta informa-
tion is reused. This may be useful when the delta information includes advertising copy
that may be repeated again.

Incorporated Disclosures

The invention described herein can be used in conjunction with inventions
described in the following applications:

- Application Serial No: 09/888,374; filed June 22, 2001, in the name of Stephane Kas-
riel, titled "Content Delivery Network using Differential Caching" attorney docket
number 155.1005.01.

1 This application is hereby incorporated by reference as if fully set forth
2 herein. They are collectively referred to as the “incorporated disclosures”.

3
4 **BRIEF DESCRIPTION OF THE DRAWINGS**

5
6 Figure 1 shows a block diagram of a system for efficient compression using
7 differential caching.

8
9 Figure 2 shows a flow diagram of a method for efficient compression using
10 differential caching.

11
12 Figure 3 shows a data flow diagram in a system for efficient compression
13 using differential caching.

14
15 Figure 4 shows a data flow diagram for efficient compression using differ-
16 ential caching using a proxy encoder server.

17
18 **Description of the Preferred Embodiment**

19
20 In the following description, a preferred embodiment of the invention is de-
21 scribed with regard to preferred process steps and data structures. Those skilled in the art
22 would recognize after perusal of this application that embodiments of the invention can

be implemented using one or more general purpose processors or special purpose processors or other circuits adapted to particular process steps and data structures described herein, and that implementation of the process steps and data structures described herein would not require undue experimentation or further invention.

Lexicography

- **Originating Server** – as used herein, an “originating server” takes on the role of a server in a client-server relationship and is the original provider of content to a client device or to a mirroring server.
- **Mirroring server** – as used herein, a “mirroring server” includes any device that takes on the role of a server in a client-server relationship and that receives requests from client devices and responds to those requests by sending content that originated (in whole or in part) at an originating server.
- **Template**– as used herein, the term “template” refers to a selected portion of a web page that is relatively unchanging. If there is no difference between different instances of a web page, then the entire page may be a template.
- **Delta information** – as used herein, the term “delta information” refers to a selected portion of a web page varies between instances of the web page.

System Elements

Figure 1 shows a block diagram of a system for efficient compression using differential caching.

A system for efficient compression using differential caching (shown by general character reference 100) includes one or more client devices 110 under the control of user 112, an originating server 120, a set of mirroring servers 130 and a communication network 140. A proxy encoder 150 may be positioned “in front” of the originating server 120.

Each client device 110 includes a processor, an input element, a presentation element, a local memory and system software. Client devices 110 further include software 114 disposed for communicating with the communication network 140 and software 116 for integrating a web page.

In a preferred embodiment, each client device 110 includes a general-purpose computer, such as a laptop or workstation. However, a client device 110 can also include (either alone or in conjunction with a laptop or workstation), a hand-held calendar (such as a “Palm Pilot” or other hand-held device), a portable computer, a special purpose computer, a cellular telephone or other telephonic device, a web server acting as the agent for a user, or another device. In alternative embodiments, a client device

1 110 may also include any other device disposed for performing all or some of the func-
2 tions described herein.

3
4 The software 114 for integrating a web page includes elements for decom-
5 pressing a template, decompressing delta information, and integrating the template and
6 delta information into a unified presentation for the user 112.

7
8 In some embodiments, the client device 110 also includes a decoder 118,
9 preferably as a browser “plug-in”. However, the decoder 118 may also be situated in
10 other locations downstream from the originating server 120, such as at a cache or firewall
11 associated with an ISP or at the edge of an enterprise network. The decoder 118 specifies
12 that it accepts delta encoding by adding information to the HTTP header. In such em-
13 bodiments, the decoder 118 may also perform the functions performed by software 114,
14 such as integrating information.

15
16 Those embodiments of the invention that do not include a decoder 118 are
17 known as “clientless”. In such instances, the server sees that the request does not come
18 from a decoder 118 and serves deltas in the form of javascript instructions and a refer-
19 ence to the template. The scripting capabilities of the client’s browser are directed to-
20 ward applying the delta to the template and displaying the HTML page. In other em-
21 bodiments, other scripting techniques may also be used.

1 The originating server 120 is under the control of an entity that provides
2 web pages for users 112. Similar to the client devices 110, the originating server 120 in-
3 cludes a processor, an input element, a presentation element, system software and a local
4 memory. However, unlike the client device 110, the originating server 120 includes a
5 cache of the complete set of compressed templates 122, a cache of delta information 124,
6 compression software 126 (for example gzip, glib or another comparable program) and
7 web server software 128. The compression software 126 is used to compress both tem-
8 plate information and delta information.

9
10 The compressed templates included in the cache of the complete set of
11 compressed templates 122 are derived from web pages that include relatively unchanging
12 elements, such as the backdrop in weather pages, blank charts used in stock pages, in-
13 complete order form pages and similar web pages with relatively unchanging elements.
14 If a web page does not include dynamic elements, the compressed template may include
15 the entire web page. In a preferred embodiment, the template is compressed and cached
16 with the complete set of compressed templates 122. It does not have to be compressed
17 again unless the information itself is updated (for example, with stylistic changes).

18
19 The delta information in the cache of delta information 124 includes those
20 portions of a set of web pages that are highly variable. This delta information can be de-
21 rived from web pages that include elements that change frequently such as a weather
22 forecast, the value of a stock, information to complete an order form and comparable in-

1 formation that is either very ephemeral or unique to a particular user 112. If the total
2 number of sets of delta information 124 associated with a particular URL is relatively low
3 (for example, a few different rotating banner advertisements) and that information is
4 relatively small, the server can maintain a table mapping uncompressed deltas to com-
5 pressed deltas. This saves computing resources because there is no need to re-compress
6 the delta information every time it is served. Optimal benefits are obtained if the delta
7 information is requested many times before the template changes.

8
9 The set of mirroring servers 130 is usually under the control of the same
10 entity that controls the originating server 120. Similar to the originating server 120, the
11 set of mirroring servers 130 includes a processor, system software, an input element, a
12 presentation element, a local memory, a cache of compressed web templates 132, com-
13 pression software 134, and web server software 136. However, unlike the originating
14 server 120, the cache of compressed web page templates 132 at the mirroring server 130
15 is not necessarily complete. Moreover, compression software 134 at the mirroring server
16 130 is generally used to compress delta information rather than templates as does the
17 compression software 126. In a preferred embodiment, the mirroring servers are rela-
18 tively more local to the client devices 110 than the originating server 120.

19
20 As noted *supra*, some embodiments of the system 100 also include a proxy
21 encoder 150. In these embodiments, the proxy encoder 150 is coupled to the originating
22 server 120, either as a separate server or as an encoder coupled to the originating server.

1 The proxy encoder server 150 compares information stored locally or at the client to
2 possibly fresher information from the mirroring server 130 or the originating server 120
3 and serves the compressed template and delta information to the client 110.
4

5 The communication network 140 is disposed for transporting compressed
6 templates, delta information and requests for web pages between the client devices 110,
7 the originating server 120 and the mirroring server 130. In a preferred embodiment, the
8 communication network 140 includes a packet switched network such as the Internet, as
9 well as (in conjunction with or instead of) an intranet, an enterprise network, an extranet,
10 a virtual private network, a virtual switched network, or a wireless network. In alterna-
11 tive embodiments, the communication network 140 may include any other set of commu-
12 nication links that couple the client devices 110 the originating server 120 and the mir-
13 roring servers 130.
14

15 *Method of Operation*

16

17 Figure 2 shows a flow diagram of a method for efficient compression using
18 differential caching.
19

20 A method 200 for efficient compression using differential caching is per-
21 formed by the system 100, including a set of client devices 110, an originating server
22 120, a set of mirroring servers 130 and a communication network 140.

1 Although described serially and in a particular sequence, in a preferred em-
2 bodiment the steps described herein can be performed concurrently or in parallel by the
3 system elements, or could be performed in a different sequence or some combination
4 thereof.

5
6 At a flow point 205, a user 112 is ready to request a web page from a mir-
7 roring server 130. In a preferred embodiment, the mirroring server 130 is relatively more
8 proximate to the user 112 than the originating server 120.

9
10 At a step 210, the user 112 causes the client device 110 to generate a re-
11 quest message to the mirroring server 130 for an identified web page. In a preferred em-
12 bodiment, each request for a web page is performed independently, even if a plurality of
13 requests are to be performed substantially simultaneously.

14
15 In other embodiments, the client device 110 may request the web page from
16 the originating server 120. Bypassing the mirroring server 120 may save bandwidth if
17 the originating server 130 is closer to the user 112 than the mirroring server 130.

18
19 When the request is made, the client device 110 indicates whether it can
20 receive delta information. If the client device 110 can receive delta information, the
21 method proceeds with step 215.

At a step 215, the mirroring server 130 determines if a compressed template for the web page is available in the cache of compressed templates 132. If necessary, the mirroring server 130 obtains the compressed template from the originating server 120. If the compressed template is unavailable from the originating server 120 (this may be the case if the template has not been requested before, or if the template has been recently changed), a template is created and compression software 126 at the originating server 120 compresses the template. The newly compressed template is sent from the originating server 120 to the mirroring server 130 where it is stored in the cache of compressed templates 132.

At a step 220, the mirroring server 130 transmits the compressed template to the client device 110.

If the mirroring server 130 were bypassed in step 210, the compressed template is sent to the client device 110 from the originating server 120.

At a step 225, the client device 110 or the mirroring server 130 (depending upon the configuration of the system 100) transmits a message to the originating server 120 for the delta information.

At a step 230, the originating server 120 identifies the delta information in the cache of delta information 124 and compresses it using the compression computer

1 program 126. Depending on the size and type of delta information 124, this compressed
2 delta information may be stored at the originating server 120 or the mirroring server 130.
3 This is particularly useful if there is a strong likelihood that the delta information will be
4 requested again.

5
6 At a step 235 the originating server 120 transmits the compressed delta in-
7 formation to the client device 110.

8
9 At step 240 the client device 110 integrates the template and delta by per-
10 forming the following substeps:

- 11
- 12 • At a substep 240(a), the template is decompressed.
- 13 • At a substep 240(b), the delta information is decompressed.
- 14 • At a substep 240(c), the decompressed template and decompressed delta are
- 15 integrated so as to form a complete web page that is presented to the user 112.
- 16

17 Given that the software 114 for integrating a web page is present at the client device 110,
18 the resources required to decompress the delta information and template and integrate
19 them do not contribute to the total resources required by the servers to compress and
20 serve the web page.

1 In other embodiments, the system 100 does not include a decoder 118.
2 These embodiments are known as “clientless”, because many of the functions normally
3 performed by the decoder 118 are performed differently. In such embodiments, the
4 originating server 120 identifies that the client can receive a delta (e.g. via a cookie), and
5 serves a delta instead of the document. The delta contains a reference to the template,
6 which can be served from either the mirroring servers 130 or from the originating server
7 120. In such embodiments, the template is compressed once and cached at the mirroring
8 server 130.

9
10 Figure 3 shows a data flow diagram in a system for efficient compression
11 using differential caching.

12
13 The system 300 includes a set of data flows for sending and receiving in-
14 formation between the client devices 110, the originating server 120 and the mirroring
15 server 130, using the communication link 140. It should be noted that it is not necessary
16 to exhaust every data flow to achieve efficient differential caching.

17
18 A data flow 310 includes messages sets between the client device 110 un-
19 der the control of the user 112 and the mirroring server 130.

20
21 Message set A from the client device 110 to the mirroring server 130 in-
22 cludes requests for web pages.

1 Message set B from the mirroring server 130 to the client device 110 in-
2 cludes the following:

- 3 • compressed templates
- 4 • compressed delta information.

5
6 A data flow 320 includes messages sets between the mirroring server 130 to
7 the originating server 120.

8
9 Message set C from the mirroring server 130 to the originating server 120
10 includes the following:

- 11 • requests for compressed templates that are not available in the cache of com-
12 pressed templates 132
- 13 • requests for delta information if the delta information is not available at the
14 mirroring server 130

15
16 Message set D from the originating server 120 to the mirroring server 130
17 includes the following:

- 18 • compressed templates
- 19 • compressed delta information

1 A data flow 330 includes messages sets between the originating server 120
2 and the client device 110.

3
4 Message set E from the client device 110 to the originating server 120 in-
5 cludes the following:

- 6 • requests for compressed templates
- 7 • requests for compressed delta information.

8
9 Message set F from the originating server 120 to the client device 110 includes
10 the following:

- 11 • compressed templates
- 12 • compressed delta information.

13
14 Figure 4 shows a data flow diagram for efficient compression using differ-
15 ential caching using a proxy encoder server.

16
17 A method 400 for efficient compression using differential caching is per-
18 formed by the system 100, including a set of client devices 110, an originating server
19 120, a set of mirroring servers 130, a proxy encoder 150 and a communication network
20 140.

1 Similar to method 200, the steps described herein can be performed concur-
2 rently or in parallel by the system elements, or could be performed in a different sequence
3 or some combination thereof.

4
5 At a flow point 405, a user 112 is ready to request a web page from a mir-
6 roring server 130. In a preferred embodiment, the mirroring server 130 is relatively more
7 proximate to the user 112 than the originating server 120.

8
9 In a step 410, the user 112 requests the web page. In a preferred embodi-
10 ment, this request is made using the decoder 118. The decoder 118 intercepts the request
11 and redirects it to the proxy encoder server 150. The decoder 118 also informs the proxy
12 encoder server 150 that the client 110 can receive delta and template information, and
13 provide delta encoding.

14
15 In other embodiments, the request can be made without using the decoder
16 118. In these other embodiments, the request goes directly to the proxy encoder server
17 150. The absence of a decoder 118 indicates that the client cannot provide delta encod-
18 ing.

19
20 In a step 415, the proxy encoder server 150 obtains the web page or a tem-
21 plate corresponding to the web page from the mirroring server or originating server.
22 The proxy encoder server 150 compares the web page or template with information in

1 the cache of compressed templates 133 (as noted *supra*, in such embodiments, these
2 caches are local to the proxy encoder server 150). If there is not a corresponding tem-
3 plate in the cache of compressed templates 126 (as may be the case if the web page was
4 never requested before), then a new template is generated, compressed and cached. If
5 there is a corresponding template in the cache of compressed templates 126 and that cor-
6 responding template needs to be updated, then updating is performed at this time and the
7 updated template is compressed and cached. The compressed template may be tagged
8 with information specifying a version number associated with the template.

9
10 In those embodiments that do not provide delta encoding, the compressed
11 template may be stored at the client 110. In such embodiments, the proxy encoder server
12 140 calculates the difference between the compressed template at the client 110 and the
13 fresh page and calculates a delta. The delta is an HTML page that includes a reference to
14 a compressed template (preferably a Javascript) and some Javascript instructions. In this
15 embodiment, the Javascript instructions tell the client 110 how to transform the template
16 into the updated web page.

17
18 In a step 420, the proxy encoder server 150 sends the compressed template
19 to the client. This compressed template may be stored at the client 110.

20
21 The following steps occur when the user 112 requests the same web page
22 associated with the compressed and cached template:

1 In a step 425, the user 112 requests a web page using the decoder 118. In
2 making the request, the decoder 118 also specifies what version of a compressed template
3 for the web page has been received in the past. This request is directed to the proxy en-
4 coder server 150.

5
6 In a step 430, the proxy encoder server 150 receives the request and identi-
7 fies a compressed template that is responsive to the request. The proxy encoder server
8 150 calculates the difference between the version of the compressed template and a pre-
9 sumably fresher version of the web page that is obtained from the originating server 120
10 or the mirroring server 130. This difference is the delta information.

11
12 In step 435, the proxy encoder sends delta information such as is responsive
13 to the request. In some embodiments, the delta information is compressed before send-
14 ing. In other embodiments, the delta information is compressed and cached.

15
16 In the “clientless” version, the proxy encoder server 150 does not know
17 what version of the template is at the client. The proxy encoder 150 makes this decision
18 and instructs the client to use a specific version of the template. This is comparable to
19 steps 310 – 335.

1 *Alternative Embodiments*

2
3 Although preferred embodiments are disclosed herein, many variations are
4 possible which remain within the concept, scope and spirit of the invention; these varia-
5 tions would be clear to those skilled in the art after perusal of this application. One such
6 variation includes caching (either at the originating server, mirroring server or proxy en-
7 coder server) a Huffman tree corresponding to a web page that has been requested in the
8 past. In such alternative embodiments, the delta is calculated by comparing the Huffman
9 tree to newer versions of the tree and computing a delta based upon those parts of the tree
10 that have changed. This technique is preferable for web pages that are requested very
11 frequently (or that change very rarely). In other embodiments, a Huffman tree corre-
12 sponding to a template is generated and served separate from the delta information.

T02030"603E2560